

Performability Analysis of Guarded-Operation Duration: A Successive Model-Translation Approach

Ann T. Tai* William H. Sanders[†] Leon Alkalai[‡] Savio N. Chau[‡] Kam S. Tso*

| | | |
|--------------------------|-------------------------------------|--|
| IA Tech, Inc.* | University of Illinois [†] | Jet Propulsion Laboratory [‡] |
| Los Angeles, CA 90024 | Urbana, IL 61801 | Pasadena, CA 91109 |
| (310) 474-3568 | (217) 333-0345 | (818) 354-3309 |
| {a.t.tai,k.tso}@ieee.org | whs@crhc.uiuc.edu | {lalkalai,schau}@jpl.nasa.gov |

Abstract

When making an engineering design decision, it is often necessary to consider its implications on both system performance and dependability. In this paper, we present a performability study that analyzes the guarded operation duration ϕ for onboard software upgrading. In particular, we define a “performability index” Y that quantifies the extent to which the guarded operation with a duration ϕ reduces the expected total performance degradation. In order to solve for Y , we propose an approach that translates a design-oriented model into an evaluation-oriented model that allows us to exploit efficient solution methods, successively closing the gap between the formulation of Y and its final solution. More specifically, we begin with constructing a design-oriented model that formulates Y and captures the collective effects on system performability of both performance overhead of guarded operation and failure behavior of the software components. We then translate this design-oriented model, through analytic manipulation, into an evaluation-oriented form that is an aggregate of constituent measures conducive to reward model solutions. Finally, based on this reward-mapping-enabled intermediate model, we specify reward structures in the composite base model which is built on three stochastic activity network (SAN) reward models. We describe the successive model-translation approach and show its feasibility for design-oriented performability modeling.

Keywords: Performability, total performance degradation, duration of guarded operation, model translation, stochastic activity networks

Submission Category: Regular paper

Acknowledgment: The work reported in this paper was supported in part by NASA Small Business Innovation Research (SBIR) Contract NAS3-99125.

Corresponding Author: Ann T. Tai, a.t.tai@ieee.org

1 Introduction

In order to protect an evolvable, distributed embedded system for long-life missions against the adverse effects of design faults introduced by an onboard software upgrade, a methodology called *guarded software upgrading* (GSU) has been developed [1, 2]. The GSU methodology is supported by a message-driven confidence-driven (MDCD) protocol that enables effective and efficient use of checkpointing and acceptance test techniques for error containment and recovery. More specifically, the MDCD protocol is responsible for ensuring that the system functions properly after a software component is replaced by an updated version during a mission, while allowing the updated component to interact freely with other components in the system. The period during which the system is under the escort of the MDCD protocol is called “guarded operation.”

Guarded operation thus permits an upgraded software component to start its service to the mission in a seamless fashion, and, in the case that the upgraded component is not sufficiently reliable and thus imposes an unacceptable risk to the mission, ensures that the system will be safely downgraded back by replacing the upgraded software component with an earlier version. It is anticipated that sensible use of this escorting process will minimize the expected total performance degradation which comprises 1) the performance penalty due to design-fault-caused failure, and 2) the performance reduction due to the overhead of the safeguard activities. Accordingly, an important design parameter is the duration of the guarded operation ϕ , as the total performance degradation is directly influenced by the length of the escorting process. In turn, this suggests that a performability analysis [3, 4] is well-suited for the engineering decision-making.

Performability analysis of this type involves some challenges. First, performability measures for engineering decision-making should be defined from a system designer’s perspective, which naturally leads to a design-oriented formulation that may not be directly conducive to the final solution. Second, such a performability model may cover a broad spectrum of system attributes with which the system designer is concerned; thus we may encounter an unmanageably large state space if we attempt to use a brute-force approach to elaborate the model.

To circumvent the difficulties, we propose an approach that solves the performability measure through successive model translation. In particular, we first define a “performability index” Y , that quantifies the extent to which the guarded operation with a duration ϕ reduces the expected total performance degradation, relative to the case in which guarded operation is completely absent. For clarity and simplicity of the design-oriented model, we allow Y to be formulated at a high level of abstraction. In order to solve for Y efficiently, we choose not to directly elaborate its formulation or map the design-oriented model to a monolithic, state-space based model. Instead, we apply analytic methods to translate the design-oriented model into an evaluation-oriented model that allows us to exploit efficient solution methods, successively closing the gap between the formulation of Y and its final solution.

More specifically, we begin with constructing a design-oriented model that formulates Y and captures the collective effects on system performability of both performance overhead of guarded operation and failure behavior of the software components. We subsequently translate this design-

oriented model, through analytic manipulation, into an evaluation-oriented form that is an aggregate of constituent measures conducive to reward model solutions. Enabled by this intermediate model, we take our final step to specify reward-structures in the composite base model, which is built on three measure-adaptive stochastic activity network (SAN) reward models.

As with hierarchical modeling techniques (see [5, 6, 7], for example) and model-decomposition methods (see [8, 9, 10], for example), the objective of this successive model-translation approach is to avoid dealing with a model that is too complex to allow derivation of a closed-form solution or has a state space that is too large to manage. The difference between those previously developed techniques and our approach is that we emphasize translating a model progressively until it reaches a form that is a simple function of “constituent reward variables,” each of which can be directly mapped to a reward structure for solution.

While the goal of the model-translation approach is clear, a complete roadmap of the translation is typically not available when the process begins. As described in Sections 3 and 4, as translation progresses, we are able to learn additional mathematical properties and implications of the behavior of the system in question. By discovering them along the path of translation, we acquire ideas for efficient model construction and solution. This is an advantage of the successive model-translation approach, because it supports performability studies of engineering problems whose mathematical properties and/or implications may not be obvious before we translate the problem formulation. Our study also shows that such a model-translation approach naturally facilitates the application of behavioral-decomposition and measure-adaptive techniques for achieving solution efficiency.

The next section provides a review of the GSU methodology and guarded operation. Section 3 defines and formulates the performability measure. Section 4 describes how we translate a design-oriented model into an evaluation-oriented model, followed by Section 5, which shows how the reward structures are specified in SAN models. Section 6 presents an analysis of optimal guarded-operation duration. The paper is concluded by Section 7, which summarizes what we have accomplished.

2 Review of Guarded Software Upgrading

The development of the GSU methodology is motivated by the challenge of guarding an embedded system against the adverse effects of design faults introduced by an onboard software upgrade [1, 2]. The performability study presented in this paper assumes that the underlying embedded system consists of three computing nodes. (This assumption is consistent with the current architecture of the Future Deliveries Testbed at JPL¹.) Since a software upgrade is normally conducted during a non-critical mission phase when the spacecraft and science functions do not require full computation power, only two processes corresponding to two different application software components are supposed to run concurrently and interact with each other. To exploit inherent system resource redundancies, we let the old version, in which we have high confidence due to its sufficiently long

¹More recently, we have extended the error containment and recovery algorithms so that the methodology can serve a more general class of distributed embedded systems [11].

onboard execution time, escort the new-version software component through two stages of GSU, namely, *onboard validation* and *guarded operation*, as illustrated in Figure 1.

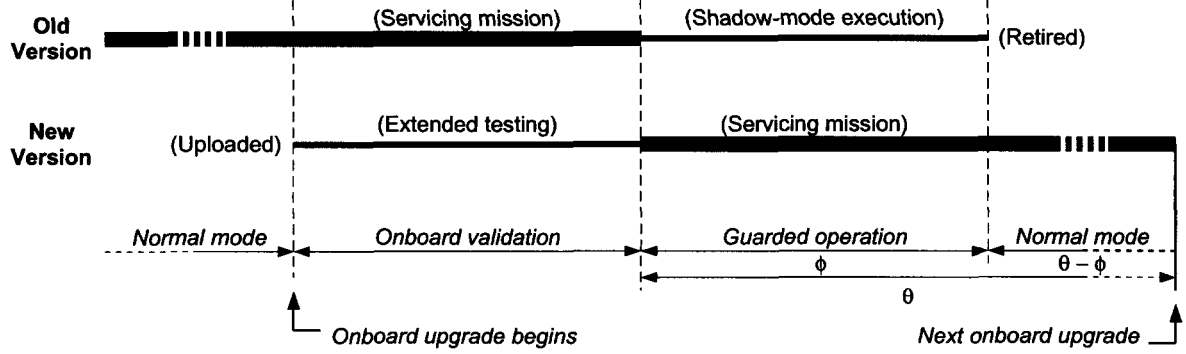


Figure 1: Onboard Guarded Software Upgrading

Further, we make use of the third processor, which would otherwise be idle during a non-critical mission phase, to accommodate the old version such that the three processes (i.e., the two corresponding to the new and old versions, and the process corresponding to the second application software component) can be executed concurrently. To aid in the description, we introduce the following notation:

- P_1^{new} The process corresponding to the new version of an application software component.
- P_1^{old} The process corresponding to the old version of the application software component.
- P_2 The process corresponding to another application software component (which is not undergoing upgrade).

The first stage of GSU (onboard validation), which can be viewed as extended testing in an actual space environment, starts right after the new version is uploaded to the spacecraft. During this stage, the outgoing messages of the shadow process P_1^{new} are suppressed but selectively logged, while P_1^{new} receives the same incoming messages that the active process P_1^{old} does. Thus, P_1^{new} and P_1^{old} can perform the same computation based on identical input data. By maintaining an onboard error log that can be downloaded to the ground for validation-results monitoring and Bayesian-statistics reliability analyses (as suggested by a number of research literatures, see [12], for example), we can make decisions regarding how long onboard validation should continue and whether P_1^{new} can be allowed to enter mission operation. If onboard validation concludes successfully, then P_1^{new} and P_1^{old} switch their roles to enter the guarded operation stage. The time to the next upgrade θ is determined upon the completion of onboard validation, according to 1) the planned duty of the flight software in the forthcoming mission phases, and 2) the quality of the flight software learned through onboard validation.

During guarded operation, P_1^{new} actually influences the external world and interacts with process P_2 under the escort of the MDCD error containment and recovery protocol, while the messages of P_1^{old} that convey its computation results to P_2 or external systems (e.g., devices) are suppressed. We call the messages sent by processes to external systems and the messages between processes *external messages* and *internal messages*, respectively.

Because the objective of the MDCD protocol is to mitigate the effect of residual software design faults, we must ensure consistency among different processes' views on verified correctness (validity) of process states and messages. Accordingly, the MDCD algorithms aim to ensure that the error recovery mechanisms can bring the system into a global state that satisfies validity-concerned global state consistency and recoverability. The key assumption used in the derivation of the MDCD algorithms is that an erroneous state of a process is likely to affect the correctness of its outgoing messages, while an erroneous message received by an application software component will result in process state contamination [13]. Accordingly, the necessary and sufficient condition for a process to establish a checkpoint is that the process receives a message that will make the process's otherwise non-contaminated state become potentially contaminated. In order to keep performance overhead low, the correctness validation mechanism, *acceptance test* (AT), is only used to validate external messages from the active processes that are potentially contaminated. By a "potentially contaminated process state," we mean 1) the process state of P_1^{new} that is created from a low-confidence software component, or 2) a process state that reflects the receipt of a not-yet-validated message that is sent by a process when its process state is potentially contaminated.

Figure 2 illustrates the behavior of the MDCD protocol. The horizontal lines in the figure represent the software executions along the time horizon. Each of the shaded regions represents an execution interval during which the state of the corresponding process is potentially contaminated. Symbols m_{ij} and M_{ik} denote, respectively, the j th internal message and k th external message sent by process P_i .




Figure 2: MDCD Protocol for Guarded Operation

Upon the detection of an erroneous external message, P_1^{old} will take over P_1^{new} 's active role and prepare to resume normal computation with P_2 . By locally checking its knowledge about whether its process state is contaminated, a process will decide to roll back or roll forward, respectively. After a rollback or roll-forward action, P_1^{old} will “re-send” the messages in its message log or further suppress messages it intends to send, based on the knowledge about the validity of P_1^{new} 's messages. After error recovery (which marks an unsuccessful but safe onboard upgrade), the system goes back to the normal mode (under which safeguard functions, namely, checkpointing and AT, are no longer performed) until the next scheduled upgrade. An undetected, erroneous external message² will result in system failure, meaning that the system will become unable to continue proper mission operation. On the other hand, as the MDCD algorithms allow error recovery to be trivial [2], we anticipate that the system will recover from an error successfully so long as the detection is successful.

If no error occurs during ϕ , then guarded operation concludes and the system goes back to normal mode at ϕ (see Figure 1). Note that while the time to the next scheduled onboard upgrade θ is chosen via a software engineering decision, the duration of guarded operation ϕ is a design parameter that influences system performance and dependability. The central purpose of this paper is to study how to evaluate a performability measure for determining an optimal ϕ . In the section that follows, we define and formulate the performability measure.

3 Performability Measure

3.1 Definition

We define a performability measure that will help us to choose the appropriate duration of guarded operation ϕ . More specifically, ϕ will be determined based on the value of the performability measure that quantifies the total performance degradation reduction resulting from guarded operation.

As mentioned in Section 1, we consider two types of performance degradation, namely:

- 1) The performance degradation caused by the performance overhead of checkpoint establishment and AT-based validation, and
- 2) The performance degradation due to design-fault-caused failure.

Clearly, a greater value of ϕ implies 1) a decrease in the performance degradation due to potential system failure caused by residual design faults in the upgraded software component, and 2) an increase in the performance degradation due to the overhead of checkpointing and AT. If we let W_ϕ denote the amount of “mission worth,” which is quantified by the system time that is devoted to performing application tasks rather than safeguard activities and accrued through θ when the duration of guarded operation (G-OP) is ϕ , then W_0 refers to the total mission worth accrued through θ for the boundary case in which the G-OP mode is completely absent (having a

²For simplicity, in the remainder of the paper, we use the term “error” to refer to an erroneous external message (to a device).

zero duration). On the other extreme, if the system is perfectly reliable, then it would not require guarded operation and would thus be free of either type of performance degradation described above. We view this extreme case as the “ideal case” and let its total mission worth (accrued through θ) be denoted by W_I .

It is worthwhile noting that the difference between the expected values of W_I and W_ϕ can be regarded as the expected mission worth reduction, or the expected total performance degradation (from the ideal case) that the system experiences through θ when the G-OP duration is ϕ . Similarly, the difference between the expected values of W_I and W_0 represents the expected total performance degradation the system experiences through θ when the G-OP mode is absent throughout θ . It follows that if $E[W_I] - E[W_\phi] < E[W_I] - E[W_0]$, then ϕ can be considered a better choice than ϕ' . Accordingly, we let the performability measure take the form of a *performability index* Y , that quantifies the extent to which a G-OP duration ϕ reduces the expected total performance degradation, relative to the case in which the G-OP mode is completely absent. More succinctly, Y is the ratio of the difference between $E[W_I]$ and $E[W_0]$ to that between $E[W_I]$ and $E[W_\phi]$:

$$Y = \frac{E[W_I] - E[W_0]}{E[W_I] - E[W_\phi]} \quad (1)$$

Based on the above discussion, we can anticipate performability benefit from a guarded operation that is characterized by a duration ϕ when $E[W_I] - E[W_\phi]$ is less than $E[W_I] - E[W_0]$. More precisely, $Y > 1$ implies that the application of guarded operation will yield performability benefit with respect to the reduction of total performance degradation. On the other hand, $Y \leq 1$ suggests that guarded operation will not be effective for total performance degradation reduction. We formulate $E[W_I]$, $E[W_0]$, and $E[W_\phi]$ in the next subsection.

3.2 Formulation

As explained above, we choose to quantify “mission worth” in terms of the system time devoted to performing application tasks (rather than safeguard activities) that is accrued through mission period $[0, \theta]$. Further, the system behavior described in Section 2 suggests that an error that propagates to an external system will nullify the worth of that mission period. Since neither of the two cases, the ideal case and the case in which the G-OP mode is completely absent, involves safeguard activities, W_I and W_0 can be formulated in a straightforward fashion:

$$W_I = 2\theta \quad (2)$$

$$W_0 = \begin{cases} 2\theta & \text{if no error occurs throughout } \theta \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

Note that the coefficient 2 in the above equations is due to the fact that in the avionics system we consider, only two application processes actively service the mission during θ . (For the cases to which W_I and W_0 correspond, the two processes will always be P_1^{new} and P_2 .)

To help to formulate an expression for W_ϕ , we group the possible behaviors (i.e., sample paths) that the system may take into several categories. In particular, since we do not make the assumption that P_1^{old} and P_2 are perfectly reliable and AT has a full coverage, we must consider situations where the system fails during guarded operation, or fails after error recovery. This leads us to define three classes of sample paths: i) those in which no error occurs, and the system thus goes through the upgrade process successfully (called S1 below), ii) those in which an error occurs during $(0, \phi]$, and the system successfully recovers (called S2 below), and iii) those involving the occurrence of an error from which the system cannot recover (no mission worth is accumulated, so these paths are not considered in the expression of mission worth). More specifically, we define sets of sample paths S1 and S2 as follows:

- S1) No error occurs by the end of ϕ , so the system enters the normal mode with P_1^{new} and P_2 in mission operation after ϕ ; the upgraded system subsequently goes through the period $(\theta - \phi)$ successfully.
- S2) An error occurs and is detected by P_1^{new} or P_2 at τ , $0 < \tau \leq \phi$, so that error recovery brings the system into the normal mode with P_1^{old} and P_2 in mission operation after τ ; the recovered system subsequently goes through $(\theta - \tau)$ successfully.

We let $\rho_{t,1}$ and $\rho_{t,2}$ denote the fractions of time during which P_1^{new} and P_2 (respectively) make forward progress (rather than performing safeguard functions), given that the system is under the G-OP mode until t ($t \leq \phi$). Then, W_ϕ can be defined as follows:

$$W_\phi = \begin{cases} (\rho_{\phi,1} + \rho_{\phi,2})\phi + 2(\theta - \phi) & \text{if system experiences a sample path in S1} \\ \gamma((\rho_{\tau,1} + \rho_{\tau,2})\tau + 2(\theta - \tau)) & \text{if system experiences a sample path in S2} \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where the coefficient γ ($0 < \gamma < 1$) is the *discount factor* that takes into account for the additional mission worth reduction for an unsuccessful but safe onboard software upgrade, relative to the case in which the upgrade succeeds. We can define γ according to the implication of S2 on the system in question. For clarity of illustration, we will postpone our description of how we define γ until Section 6, in which we present the evaluation experiments and results.

In order to solve for Y , we first define a stochastic process, $\mathcal{X} = \{X_t \mid t \in [0, \theta]\}$, to represent the dynamics of the distributed embedded system that is undergoing an onboard upgrade. As mentioned in Section 2, when an error goes undetected, the system will lose its ability to continue mission operation, implying an absorbing state. If we let \mathcal{A}_1 denote the set of states of \mathcal{X} in which no error has occurred in the system, then according to the definition of W_0 ,

$$E[W_0] = 2\theta P(X_\theta \in \mathcal{A}_1, \text{ when G-OP duration is } 0). \quad (5)$$

Further, if we let $W_\phi^{\text{S1}} = (\rho_{\phi,1} + \rho_{\phi,2})\phi + 2(\theta - \phi)$ and $W_\phi^{\text{S2}} = \gamma((\rho_{\tau,1} + \rho_{\tau,2})\tau + 2(\theta - \tau))$ (see Equation (4)), then by the theorem of total expectation,

$$E[W_\phi] = E[W_\phi^{S1}] + E[W_\phi^{S2}]. \quad (6)$$

By definition, the stochastic process \mathcal{X} will be in a state in \mathcal{A}_1 at θ if the system has a G-OP duration ϕ and experiences a sample path in S1. It follows that

$$E[W_\phi^{S1}] = ((\rho_{\phi,1} + \rho_{\phi,2})\phi + 2(\theta - \phi)) P(X_\theta \in \mathcal{A}_1, \text{ when G-OP duration is } \phi) \quad (7)$$

We notice that the application-purpose message-passing events that trigger checkpointing and AT (which dominate the performance overhead) are significantly more frequent than the fault-manifestation events. Moreover, the mean time between message-passing events is only seconds in length, whereas a reasonable value of ϕ will be in the range of hundreds or thousands of hours. Hence, we can assume that the system reaches a steady state with respect to the performance-overhead related events before an error occurs or the G-OP duration ends. Thus, $\rho_{t,1}$ and $\rho_{t,2}$ can be regarded as steady-state measures ρ_1 and ρ_2 , respectively. Consequently, Equation (7) becomes:

$$E[W_\phi^{S1}] = ((\rho_1 + \rho_2)\phi + 2(\theta - \phi)) P(X_\theta \in \mathcal{A}_1, \text{ when G-OP duration is } \phi) \quad (8)$$

Although combinatoric-form expressions can be derived at the design-oriented level for $E[W_0]$ and $E[W_\phi^{S1}]$, the complexity of sample paths in S2, coupled with the fact that τ is a random variable that can assume a continuum of values, precludes the possibility of deriving a combinatoric-form expression for $E[W_\phi^{S2}]$. Accordingly, we let h be the probability density function (pdf) of τ and f denote the pdf of the time to system failure that occurs after error recovery (when P_1^{old} and P_2 are in the mission operation). Then, $E[W_\phi^{S2}]$ can be formulated as follows:

$$E[W_\phi^{S2}] = \gamma \int_0^\phi ((\rho_1 + \rho_2)\tau + 2(\theta - \tau))h(\tau) \left(1 - \int_\tau^\theta f(x) dx\right) d\tau \quad (9)$$

To this end, we arrive at a stage in which Y is expressed in “conceptual mathematical terms” that are design-oriented and require further elaboration. In particular, most of the terms in Equations (5), (8), and (9), are not ready to be solved for directly. If we chose a “non-state-space approach” and began with elaborating h and f , then Equation (9) would involve a very complicated convolution that we would be unable to solve without excessive approximation. Whereas if we wish to solve those equations using reward model solution techniques and attempt to build a monolithic model for this purpose, there would be no guarantee that the resulting model could support the evaluation of those equations; even if support is available when the model is constructed in a comprehensive way, the complexity of the model may prevent us from obtaining numerical results efficiently. Those factors collectively suggest that rather than take an approach that deals with a complex model using numerical algorithms, we should endeavor to avoid or reduce model complexity by translating the design-oriented model into a form for which we are able to exploit techniques such as behavioral decomposition and hierarchical composition of reward-model solutions.

4 Successive Model-Translation Approach

4.1 Translation toward Reward Model Solutions

With the motivation described at the end of the previous section, we develop an approach that translates the design-oriented model step by step until it reaches a stage at which the final solution of Y becomes a simple function of “constituent measures,” each of which can be directly mapped to a reward structure. Figure 3 illustrates the process of successive model translation.

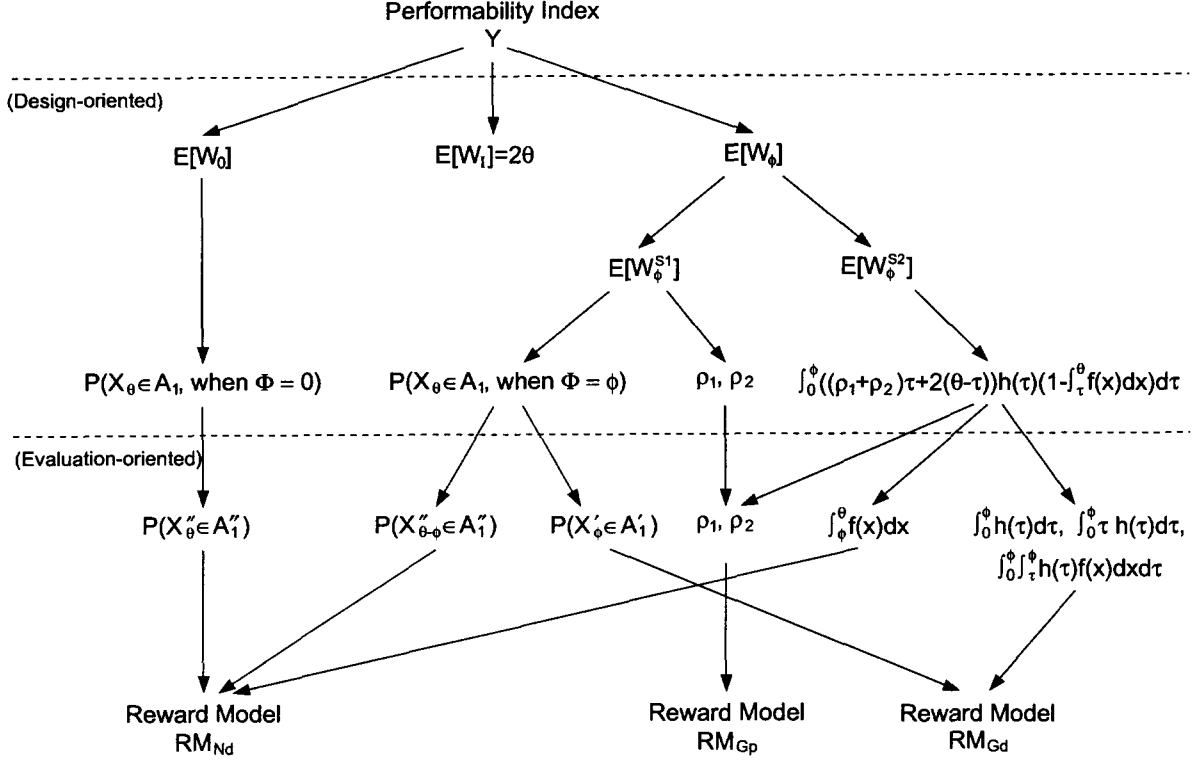


Figure 3: Successive Model Translation

As described in Section 3.2 and shown in Figure 3, the design-oriented formulation of Y results in the following mathematical terms:

- $P(X_\theta \in A_1, \text{ when G-OP duration is } 0)$ in Equation (5).
- $P(X_\theta \in A_1, \text{ when G-OP duration is } \phi)$, in Equation (8).
- ρ_1, ρ_2 in Equations (8) and (9).
- The double integral in Equation (9), i.e., $\int_0^\phi ((\rho_1 + \rho_2)\tau + 2(\theta - \tau))h(\tau) \left(1 - \int_\tau^\theta f(x) dx\right) d\tau$.

Since these terms are at a high level of abstraction and cannot be solved directly, we perform translation by applying analytic techniques to realize model decomposition and measure partition/conversion. It is possible to specify a monolithic model to represent the stochastic process \mathcal{X}

for solving $P(X_\theta \in \mathcal{A}_1, \text{when G-OP duration is } 0)$ and $P(X_\theta \in \mathcal{A}_1, \text{when G-OP duration is } \phi)$, if we choose to use a model type that is highly expressive, such as stochastic Petri nets or stochastic activity networks. However, as mentioned previously, the complexity of the model would make it impossible for us to achieve solution efficiency, even if the model is comprehensive enough to support the measure. Therefore, we let the stochastic process \mathcal{X} be partitioned into two simpler processes, namely, $\mathcal{X}' = \{X'_t \mid t \in [0, \phi]\}$ and $\mathcal{X}'' = \{X''_t \mid t \in [0, \theta]\}$. The former represents the system behavior during the pre-designated G-OP interval³. The latter can represent the system behavior under the normal mode in different situations: 1) after G-OP completes successfully, and 2) when the G-OP mode is completely absent during $[0, \theta]$. Furthermore, with a reasonably high message-sending rate, the likelihood that dormant error conditions will remain in a process state after error recovery is so low that the effect on system behavior is practically negligible [13]. Therefore, \mathcal{X}'' can also represent the system behavior after the pre-designated G-OP interval and provide a good approximation for dependability measures, for the case in which an error occurs and is detected in the system when it is under the G-OP mode.

Based on the decomposition, we are able to solve for $P(X_\theta \in \mathcal{A}_1, \text{when G-OP duration is } 0)$ and $P(X_\theta \in \mathcal{A}_1, \text{when G-OP duration is } \phi)$ in a way that is significantly more efficient. More specifically, the former can be converted into $P(X''_\theta \in \mathcal{A}''_1)$ while the latter can be translated as the product of $P(X'_\phi \in \mathcal{A}'_1)$ and $P(X''_{\theta-\phi} \in \mathcal{A}''_1)$, if we let \mathcal{A}'_1 and \mathcal{A}''_1 denote, respectively, the sets of states of \mathcal{X}' and \mathcal{X}'' in which no error has occurred in the system. Consequently, each of those transient, instant-of-time measures can be solved by defining a reward structure in one of the decomposed models, as illustrated in Figure 3.

As explained in Section 3.2, we treat ρ_1 and ρ_2 as steady-state instant-of-time measures. This suggests that we can evaluate these two constituent measures in a reward model that represents the performance aspects of the stochastic process \mathcal{X}' (and has no absorbing states). In other words, as illustrated in Figure 3, ρ_1 and ρ_2 are ready for reward model solutions, requiring no further translation.

Clearly, it is more challenging to translate the double integral in Equation (9) into a form that is conducive to a reward model solution. Hence, we use judgment and make decisions regarding how to proceed along the path of translation. Step by step, as described in detail in Section 4.2, the double integral is converted into a form that is an aggregate of constituent measures, namely, ρ_1 , ρ_2 , $\int_0^\phi h(\tau) d\tau$, $\int_0^\phi \tau h(\tau) d\tau$, $\int_0^\phi \int_\tau^\phi h(\tau) f(x) dx d\tau$, and $\int_\phi^\theta f(x) dx$. For each of those measures, we can define a reward structure in one of the decomposed models that represents \mathcal{X}' or \mathcal{X}'' and supports dependability or performance-overhead measures (as described in the preceding paragraphs). By solving the above constituent reward variables individually, we can evaluate the translated form of the double integral in an efficient fashion.

In summary, the translation process described thus far converts the terms at the design-oriented level into the following solvable constituent reward variables (see also Figure 3):

³The behavior of a recovered system within the interval $(\tau, \phi]$ can also be represented by \mathcal{X}' , if an error occurs and is detected at τ ($\tau \leq \phi$).

- $P(X''_{\theta} \in \mathcal{A}_1)$, $P(X'_{\phi} \in \mathcal{A}_1)$, and $P(X''_{\theta-\phi} \in \mathcal{A}_1)$.
- ρ_1, ρ_2 .
- $\int_0^{\phi} h(\tau) d\tau$, $\int_0^{\phi} \tau h(\tau) d\tau$, $\int_0^{\phi} \int_{\tau}^{\phi} h(\tau) f(x) dx d\tau$, and $\int_{\phi}^{\theta} f(x) dx$.

To this end, it becomes apparent that we will be able to solve for Y if we construct the following three reward models at the base-model level:

RM_{Gd} A reward model that represents the system behavior during the pre-designated G-OP interval and supports dependability measures.

RM_{Nd} A reward model that represents the system behavior under the normal mode and supports dependability measures.

RM_{Gp} A reward model that represents the system behavior under the G-OP mode and supports performance-overhead measures.

Details about the mapping between the resulting constituent measures and the reward structures in RM_{Gd} , RM_{Nd} , and RM_{Gp} are provided in Section 5.

4.2 Translation of $E[W_{\phi}^{S2}]$

While we have explained in a fairly detailed way other aspects of the translation process, we now provide a detailed description of the successive translation of the double integral in the expression for $E[W_{\phi}^{S2}]$ (see Equation (9)). We begin with rearranging its terms:

$$E[W_{\phi}^{S2}] = \gamma \left(\int_0^{\phi} (2\theta - (2 - (\rho_1 + \rho_2))\tau) h(\tau) d\tau - \int_0^{\phi} (2\theta - (2 - (\rho_1 + \rho_2))\tau) h(\tau) \int_{\tau}^{\theta} f(x) dx d\tau \right) \quad (10)$$

If further, we rearrange the first term in the parentheses of Equation (10), we have

$$\int_0^{\phi} (2\theta - (2 - (\rho_1 + \rho_2))\tau) h(\tau) d\tau = 2\theta \int_0^{\phi} h(\tau) d\tau - (2 - (\rho_1 + \rho_2)) \int_0^{\phi} \tau h(\tau) d\tau \quad (11)$$

Note that τ has a mixture distribution [14]. This is because $h(\tau)$ equals zero for $\tau > \phi$ and thus $\lim_{\tau \rightarrow \infty} H(\tau) < 1$. Clearly, $\int_0^{\phi} h(\tau) d\tau$ is the probability that an error occurs and is detected by ϕ when the G-OP duration is ϕ . However, as mentioned earlier, the complexity of the system behavior, makes it very difficult to derive h and compute the integrals without an excessive amount of approximation. Therefore, we choose to use reward model solution techniques and assume that rewards are associated with the states of \mathcal{X}' . More specifically, we let \mathcal{A}'_3 denote the set of states (of \mathcal{X}') in which an error has occurred and been successfully detected, $\int_0^{\phi} h(\tau) d\tau$ can then be evaluated

as the expected instant-of-time reward:

$$\int_0^\phi h(\tau) d\tau = P(X'_\phi \in \mathcal{A}'_3) \quad (12)$$

In other words, with a state-space based model \mathcal{X}' , we can solve $\int_0^\phi h(\tau) d\tau$ by assigning a reward rate of 1 to all states in \mathcal{A}'_3 and a reward rate of zero to all other states, and computing the expected reward at ϕ .

Recall that the system behavior implies that 1) a state in which the system encounters an undetected error is absorbing, and 2) a successful error detection will result in error recovery that brings the system back to the normal mode under which checkpointing and AT (the error detection mechanism) will no longer be performed. In turn, this suggests that mean time to error detection $\int_0^\phi \tau h(\tau) d\tau$ is a meaningful measure, and it can have a reward model solution. Accordingly, if we let \mathcal{A}'_2 denote the set of states in which no error has been detected, and \mathcal{A}'_4 denote the set of (absorbing) states in which an error has occurred and caused a system failure due to unsuccessful error detection (thus \mathcal{A}'_4 is a proper subset of \mathcal{A}'_2), we have

$$\int_0^\phi \tau h(\tau) d\tau = \int_0^\phi (P(X'_t \in \mathcal{A}'_2) - P(X'_t \in \mathcal{A}'_4)) dt \quad (13)$$

which implies that, to solve $\int_0^\phi \tau h(\tau) d\tau$, we can assign a reward rate of 1 to all states (of \mathcal{X}') in \mathcal{A}'_2 , a reward rate of -1 to all states in \mathcal{A}'_4 , and a reward rate of zero to all other states, and then compute the expected reward accumulated through ϕ . Thus, the integrals in Equation (11) (and thus the first term of Equation (10)) can be solved.

We manipulate the second term in Equation (10) in a similar fashion and begin with rearranging the terms:

$$\begin{aligned} & \int_0^\phi (2\theta - (2 - (\rho_1 + \rho_2))\tau) h(\tau) \int_\tau^\theta f(x) dx d\tau \\ &= 2\theta \int_0^\phi \int_\tau^\theta h(\tau) f(x) dx d\tau - (2 - (\rho_1 + \rho_2)) \int_0^\phi \int_\tau^\theta \tau h(\tau) f(x) dx d\tau \\ &\approx 2\theta \int_0^\phi \int_\tau^\theta h(\tau) f(x) dx d\tau \end{aligned} \quad (14)$$

We neglect the term $(2 - (\rho_1 + \rho_2)) \int_0^\phi \int_\tau^\theta \tau h(\tau) f(x) dx d\tau$ because its value differs from those of θ and $E[W_\phi^{S2}]$ by orders of magnitude. Note also that the expression $\int_0^\phi \int_\tau^\theta h(\tau) f(x) dx d\tau$ is the probability that an error occurs and is detected when the system is under the G-OP mode but the system fails due to another error that occurs between the successful recovery and the next upgrade. Nonetheless, the area of the integration goes across the boundary between the pre-designated G-OP interval $[0, \phi]$ and the interval $[\phi, \theta]$ during which the system (that completes G-OP safely) continues mission operation under the normal mode. This prevents us from obtaining a reward model solution based on the decomposed submodels RM_{Gd} and RM_{Nd} . By closely inspecting the

area of the integration, as shown in Figure 4(a), we recognize that we can change the coordinates of the integrals such that the orientation of the integration area will be converted accordingly, as shown in Figure 4(b). In turn, the converted integration area suggests that the result of Equation (14) can be broken down into two terms:

$$\begin{aligned}
& 2\theta \int_0^\phi \int_\tau^\theta h(\tau) f(x) dx d\tau \\
&= 2\theta \int_0^\phi \int_0^x f(x) h(\tau) d\tau dx + 2\theta \int_\phi^\theta f(x) \int_0^\phi h(\tau) d\tau dx
\end{aligned} \tag{15}$$




Figure 4: Translating the Area of Integration

The second term of Equation (15) can thus be expressed as the product of two probabilities, namely, $\int_0^\phi h(\tau) d\tau$ and $\int_\phi^\theta f(x) dx$ (see Equation (16)). While we have already provided a reward model solution for the former (see Equation (12)), we recognize that the latter is the probability that the recovered system will fail due to the occurrence of another error at a time instant in $[\phi, \theta]$. As explained in Section 4.1, we can obtain a good approximation for $\int_\phi^\theta f(x) dx$ by defining a reward structure in RM_{Nd} (which represents the dependability aspects of \mathcal{X}'') and computing the expected instant-of-time reward at $(\theta - \phi)$.

On the other hand, the first term in Equation (15) is not in a form that can be interpreted in a straightforward way. We therefore “change back” the coordinates for this individual constituent measure so that its area of integration is translated from the darker region (with a triangular shape) in Figure 4(b) back to the darker region in Figure 4(a). Thus, the second term of Equation (10) finally becomes:

$$\begin{aligned}
& \int_0^\phi (2\theta - (2 - (\rho_1 + \rho_2)) \tau) h(\tau) \int_\tau^\theta f(x) dx d\tau \\
&= 2\theta \int_0^\phi \int_\tau^\phi h(\tau) f(x) dx d\tau + 2\theta \left(\int_0^\phi h(\tau) d\tau \right) \left(\int_\phi^\theta f(x) dx \right)
\end{aligned} \tag{16}$$

The first term (in the above equation) can now be interpreted in a straightforward manner as the probability that an error is detected when the system is under the G-OP mode and the recovered system fails by ϕ (under the normal mode) due to the occurrence of another error. A reward structure can then be defined accordingly in the reward model RM_{Gd} (which represents the dependability aspects of \mathcal{X}'). To this end, we can evaluate each of the constituent measures of $E[W_\phi^{S2}]$ by mapping it to a reward structure in RM_{Gd} , RM_{Nd} , or RM_{Gp} . In other words, if we plug the results of Equations (11) and (16) into Equation (10), $E[W_\phi^{S2}]$ becomes ready to be solved using reward model solution techniques.

5 SAN Reward Model Solutions for Constituent Measures

We use stochastic activity networks to realize the final step of model-translation. This choice is based on the following factors: 1) SANs have high-level language constructs that facilitate marking-dependent model specifications and representation of dependencies among system attributes, and 2) the *UltraSAN* tool provides convenient specification capability for defining reward structures [15], and 3) by adopting and making necessary modifications to the SAN models we developed for our previous (separate) dependability and performance-cost studies [13, 2], we are able to use them as the reward models RM_{Gd} , RM_{Nd} , and RM_{Gp} . In the following subsections, we review these SAN models briefly and describe how the SAN-based reward structures are implemented.

5.1 SAN Reward Models

The rich syntax and marking-dependent specification capability of SANs allow us to specify every aspect of the protocol precisely. However, we may encounter a state-space explosion or experience very slow computation if we attempt to construct a monolithic model, or attempt to make a SAN model a procedural specification of the MDCD protocol. To avoid these problems, we

- 1) Use three separate reward models, each of which is specified for representing the dependability or performance-overhead related aspects of \mathcal{X}' or \mathcal{X}'' , as explained in Section 4.1, and
- 2) Minimize explicit representation of the algorithmic details, while ensuring that every aspect of their impact on the particular measure we seek to solve (in a particular reward model) is captured.

The SAN reward model RM_{Gd} is a modified version of the model we built for studying the dependability gain from the use of the MDCD protocol [13]. Modifications are made so that whether an error has been detected in the system is explicitly represented; and thus each of the constituent measures $\int_0^\phi h(\tau) d\tau$, $\int_0^\phi \tau h(\tau) d\tau$, and $\int_0^\phi \int_\tau^\phi h(\tau) f(x) dx d\tau$ can be easily mapped to a reward structure. In model construction, we avoid modeling details about checkpoint establishment, deletion, and rollback error recovery. Rather, by exploiting the relations among the markings of the places that represent whether a process is actually error-contaminated and the process's knowledge

about its state contamination, we are able to characterize the system’s failure behavior precisely with respect to whether messages sent by potentially contaminated processes will cause system failure. The detailed description of the model is omitted here but can be found in [13].

In contrast, in the SAN reward model RM_{Gp} (see Figure 5), we omit those failure-behavior-related aspects, such as error occurrence and unsuccessful error detection [2]. Instead, we focus on representing those conditions that would require a process to take actions that do not belong to the category of “forward progress” (e.g., the action to establish a checkpoint, or to perform an AT).

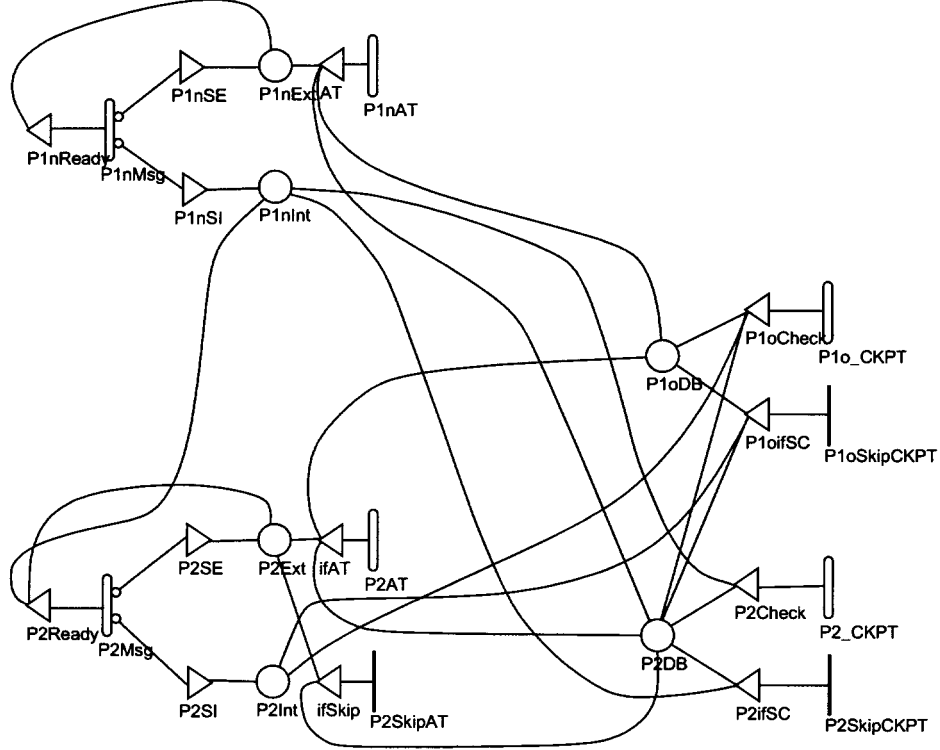


Figure 5: SAN Reward Model for Solving Performance-Overhead Measures (RM_{Gp})

Since its purpose is to solve ρ_1 and ρ_2 , this reward model includes the timed and instantaneous activities that represent the error containment actions of the protocol that are driven by the message passing events and dynamically adjusted confidence in processes, as shown in Figure 5. In particular, the places $P1oDB$ and $P2DB$ represent the dirty bits of P_1^{old} and P_2 , respectively. Each of those places may have a marking of zero or one, which can be interpreted as the confidence in the corresponding process. The timed and instantaneous activities together precisely represent how the MDCD checkpointing rule and AT-based validation policy are executed. For example, when being activated, the timed activity $P2_CKPT$ indicates that P_2 is engaged in a checkpoint establishment, while the instantaneous activity $P2SskipCKPT$ indicates that P_2 is exempted from checkpointing for a particular message receiving event, according to the MDCD checkpointing rule.

It is also worth noting that, due to the nature of the measure, the SAN reward model RM_{Gd} emphasizes the effects on system dependability of the interactions between the non-ideal environment conditions and the behavior of the MDCD protocol. Accordingly, in the model construction, we relax the design assumptions for an ideal execution environment. In contrast, the purpose of RM_{Gp} is to evaluate the performance overhead resulting from processes' error containment activities. Since those fault tolerance mechanisms are directly influenced by the design assumptions, the ideal environment assumptions are preserved in this SAN reward model. The SAN model RM_{Nd} is rather simple and thus is not described here.

5.2 Reward Model Solutions

As a result of model translation, each of the constituent measures becomes in a form that can be easily solved by defining a reward structure in the corresponding SAN reward model. In addition, the *UltraSAN* tool provides us with a convenient way to define a reward structure by specifying a "predicate-rate" pair [15]. Below we describe how the reward structures are specified in each of the SAN reward models.

5.2.1 Solving Constituent Measures in RM_{Nd}

As indicated in Figure 3 and explained in Section 4.1, three constituent measures are supposed to be solved in the reward model RM_{Nd} , namely, $P(X''_{\theta} \in \mathcal{A}''_1)$, $P(X''_{\theta-\phi} \in \mathcal{A}''_1)$, and $\int_{\phi}^{\theta} f(x) dx$. In particular, these three measures can share the following predicate-rate pair:

- *Predicate:* MARK(failure) == 0
- *Rate:* 1

To solve $P(X''_{\theta} \in \mathcal{A}''_1)$ and $P(X''_{\theta-\phi} \in \mathcal{A}''_1)$, we assign the fault-manifestation rate of P_1^{new} to the activity that represents the fault-manifestation behavior of the first software component, and compute the expected reward values at θ and $(\theta - \phi)$, respectively. As to $\int_{\phi}^{\theta} f(x) dx$, since it can be treated as the probability that the recovered system (consisting of P_1^{old} and P_2) fails during the interval $[0, \theta - \phi]$, we assign the fault-manifestation rate of P_1^{old} to the activity that represents the fault-manifestation behavior of the first software component, and compute the complement of the expected reward value at $(\theta - \phi)$.

5.2.2 Solving Constituent Measures in RM_{Gd}

As indicated in Figure 3, the constituent measures $\int_0^{\phi} h(\tau) d\tau$, $\int_0^{\phi} \tau h(\tau) d\tau$, $\int_0^{\phi} \int_{\tau}^{\phi} h(\tau) f(x) dx d\tau$, and $P(X'_{\phi} \in \mathcal{A}'_1)$ are supposed to be evaluated in the reward model RM_{Gd} . Table 1 summarizes how reward structures are specified and how the expected reward values are computed for solving those measures. An explanation of those reward model solutions has been given in Section 4.

Table 1: Constituent Measures and SAN Reward Structures in RM_{Gd}

| Measure | Reward Type | Predicate-Rate Pair | |
|--|--|--|----|
| $\int_0^\phi h(\tau) d\tau$ | Expected instant-of-time reward at ϕ | MARK(detected)==1 && MARK(failure)==0 | 1 |
| $\int_0^\phi \tau h(\tau) d\tau$ | Expected accumulated interval-of-time reward for $[0, \phi]$ | MARK(detected)==0 | 1 |
| | | MARK(detected)==0 && MARK(failure)==1 | -1 |
| $\int_0^\phi \int_\tau^\phi h(\tau) f(x) dx d\tau$ | Expected instant-of-time reward at ϕ | MARK(detected)==1 && MARK(failure)==1 | 1 |
| $P(X'_\phi \in \mathcal{A}'_1)$ | Expected instant-of-time reward at ϕ | MARK(detected)==0 && MARK(failure)==0 | 1 |

5.2.3 Solving Constituent Measures in RM_{Gp}

As indicated in Figure 3, two constituent measures are supposed to be solved in the reward model RM_{Gp} , namely, ρ_1 and ρ_2 . For simplicity and clarity of the specification of the predicate-rate pairs, we instead solve for $(1 - \rho_1)$ and $(1 - \rho_2)$, which are the performance overhead of P_1^{new} and P_2 , respectively. Table 2 enumerates the reward type and predicate-rate pair for each of the two measures.

Table 2: Constituent Measures and SAN Reward Structures in RM_{Gp}

| Measure | Reward Type | Predicate-Rate Pair | |
|--------------|---|--|---|
| $1 - \rho_1$ | Expected instant-of-time reward at steady state | MARK(P1nExt)==1 | 1 |
| $1 - \rho_2$ | Expected instant-of-time reward at steady state | (MARK(P1nInt)==1 && MARK(P2DB) == 0) (MARK(P2Ext)==1 && MARK(P2DB) == 1) | 1 |

Note that the predicate-rate pair specified for $(1 - \rho_2)$ involves more conditions. This is because, unlike the process P_1^{new} which is always considered as potentially contaminated when the system is under the G-OP mode, we dynamically adjust the confidence in P_2 and perform checkpointing and AT accordingly.

6 Evaluation Results

Applying the SAN reward models described in Section 5 and *UltraSAN*, we evaluate the performance index Y . Before we proceed to discuss the numerical results, we define the following notation:

| | |
|--------------------|---|
| μ_{new} | Fault-manifestation rate of the process corresponding to the newly upgraded software version. |
| μ_{old} | Fault-manifestation rate of a process corresponding to an old software version. |
| c | Coverage of an acceptance test. |
| λ | Message-sending rate of a process. |
| p_{ext} | Probability that the message a process intends to send is an external message. |
| α | Acceptance-test completion rate. |
| β | Checkpoint-establishment completion rate. |

We begin with conducting a study of the optimality of the G-OP duration ϕ , considering the impact of the fault-manifestation rate of the upgraded software component. Specifically, we use the parameter values shown in Table 3, in which all the parameters involving time presume that time is quantified in hours. Accordingly, $\lambda = 1200$ means that the time between message sending events (for an individual process) is 3 seconds; similarly, $\alpha = 6000$ and $\beta = 6000$ imply that the mean time to the completion of an AT-based validation and the mean time to the completion of a checkpoint establishment are both 600 milliseconds. Further, we let the γ (see Equation (4)) be a decreasing function of $\bar{\tau}$, the mean time to error detection. More succinctly, $\gamma = 1 - \frac{\bar{\tau}}{\theta}$. This function is defined based on the following consideration. Safeguard activities would no longer be performed after τ when error detection brings the system back to the normal mode with P_1^{old} and P_2 in mission operation; since that implies an unsuccessful (but safe) onboard upgrade, the performance cost paid for the safeguard activities up to τ would yield an additional reduction of mission worth, relative to the case of a successful onboard upgrade.

Table 3: Parameter Value Assignment

| θ | λ | μ_{new} | μ_{old} | c | P_{ext} | α | β |
|----------|-----------|--------------------|--------------------|------|------------------|----------|---------|
| 10000 | 1200 | 10^{-4} | 10^{-8} | 0.95 | 0.1 | 6000 | 6000 |

The numerical results from this study are displayed by the curve with solid dots in Figure 6. The values of the performability index indicate that the optimal duration of the G-OP mode for this particular setting is 7000 hours, which yields the best worth of the mission period θ , due to the greatest possible reduction of expected total performance degradation. This implies that for this particular setting, a ϕ smaller than 7000 would lead to a greater expected performance degradation due to increased risk of potential design-fault-caused failure. On the other hand, if we let ϕ be larger, then the increased performance degradation due to performance overhead would more than negate the benefit from the extended guarded operation.

By decrementing the fault-manifestation rate of P_1^{new} (μ_{new}) to 0.5×10^{-4} (while letting other parameter values remain the same), we obtain another set of values of Y , as illustrated by the companion curve marked by hollow dots in Figure 6. The two curves together reveal that the optimality of ϕ is very sensitive to the reliability of the upgraded software component. In particular,

we observe that when μ_{new} is decremented from 10^{-4} to 0.5×10^{-4} , the optimal ϕ is dropped from 7000 to 5000 (hours), even though the performance costs of safeguard activities remain low (thus ρ_1 and ρ_2 remain high, and equal 0.98 and 0.95, respectively). While it is quite obvious that a smaller μ_{new} will favor a shorter duration of the G-OP mode, this study confirms the relation between the two system attributes and helps us to recognize the sensitivity of this relation.




Figure 6: Effect of Fault-Manifestation Rate on Optimal G-OP Duration ($\theta = 10000$)

In the next study, we change the values of α and β to 2500 (i.e., the times to completions of an AT-based validation and a checkpoint establishment become 1440 milliseconds, up from 600 milliseconds in the previous study), implying that the performance costs for safeguard activities become higher. The evaluation results are shown by the curve with tiny hollow triangles in Figure 7, in which we duplicate the curve with solid dots from Figure 6 for comparison. With the decremented values of α and β , ρ_1 and ρ_2 are reduced to 0.95 and 0.90, respectively. As shown by the curve with the tiny hollow triangles, the optimal ϕ for this case is 6000, down from 7000. The change of the optimal ϕ is again a result of the tradeoffs between the two types of expected performance degradation. More specifically, this change is due to the factor that the increased performance overhead tends to further negate dependability benefits, and thus suggests an earlier cutoff line for guarded operation.

Note that so far we have used θ , the time to the next upgrade, as a constant. However, as described in Section 2, θ is indeed chosen based on a software engineering decision (at the time onboard validation completes); the decision depends upon at least two factors: 1) the planned duty of the flight software in the forthcoming mission phases, and 2) the quality of the flight software learned through onboard validation. Accordingly, we analyze the effects of the value of θ on the




Figure 7: Effect of Performance Overhead on Optimal G-OP Duration ($\theta = 10000$)

optimality of ϕ . Specifically, we repeat the study that yields the results shown in Figure 6, by letting θ be reduced to 5000 hours. The resulting curves are displayed in Figure 8.




Figure 8: Effect of Fault-Manifestation Rate on Optimal G-OP Duration ($\theta = 5000$)

It is interesting to observe that, while other parameter values remain the same (meaning that the performance and dependability attributes of the system itself do not differ from the previous study), the reduction of θ significantly changes the values for the optimal ϕ . Specifically, the optimal values of ϕ for the cases in which μ_{new} equals 10^{-4} and 0.5×10^{-4} go down to 2500 and 2000, respectively. This can be understood by considering that reliability is generally a decreasing function of time, for a system without maintenance. More precisely, when the anticipated time to the next upgrade becomes shorter, the likelihood that the system will fail before the forthcoming upgrade activity decreases, permitting guarded operation to end at an earlier point to minimize the expected total performance degradation. By inspecting the results of the constituent measures that are available to us, namely, $P(X''_{\theta} \in \mathcal{A}''_1)$ and $\int_0^{\phi} h(\tau) d\tau$, we are able to validate this explanation.

7 Concluding Remarks

We have conducted a model-based performability study that analyzes the guarded-operation duration for onboard software upgrading. By translating a design-oriented model into an evaluation-oriented model, we are able to reach a reward model solution for the performability index Y that supports the decision on the duration of guarded operation.

The model-translation approach offers us flexibility. In particular, the approach suggests to us that model translation can be done within the same level or across different levels of abstraction, depending upon the application. In addition, since its goal is to transform the problem of solving a performability measure into that of evaluating constituent reward variables, the successive model-translation approach naturally enables us 1) to explore the *imbedded* possibilities of applying efficient model construction and solution strategies, including reward model solution techniques, behavioral decomposition, measure-adaptive model construction, and hierarchical composition; and 2) to access the results of the constituent measures to gain more insights from a model-based performability evaluation.

Moreover, as noted earlier and exemplified by our performability study, we make decisions on how to further translate a model based on intermediate results obtained during successive model translation, as opposed to having a complete picture of the hierarchy prior to model translation. In other words, although our model-translation approach can be viewed as a top-down process, judgment and decisions on how to proceed are made in an “on-the-fly” fashion. This is an advantage of the successive model-translation approach, because it enables us to solve engineering problems whose mathematical properties and/or implications may not become apparent until we elaborate the formulation of the problem to a certain degree. Our current effort is directed toward continuing this investigation by carrying out more case studies in further depth. In addition, as we have developed the GSU middleware and have been in the process of porting it to the Future Deliveries Testbed at JPL, we intend to experimentally validate the parameter values used in our analysis and the results of the constituent measures.

References

- [1] A. T. Tai, K. S. Tso, L. Alkalai, S. N. Chau, and W. H. Sanders, "On low-cost error containment and recovery methods for guarded software upgrading," in *Proceedings of the 20th International Conference on Distributed Computing Systems (ICDCS 2000)*, (Taipei, Taiwan), pp. 548–555, Apr. 2000.
- [2] A. T. Tai, K. S. Tso, L. Alkalai, S. N. Chau, and W. H. Sanders, "Low-cost error containment and recovery for onboard guarded software upgrading and beyond," *IEEE Trans. Computers*, vol. 51, Feb. 2002.
- [3] J. F. Meyer, "On evaluating the performability of degradable computing systems," *IEEE Trans. Computers*, vol. C-29, pp. 720–731, Aug. 1980.
- [4] J. F. Meyer, "Performability: A retrospective and some pointers to the future," *Performance Evaluation*, vol. 14, pp. 139–156, Feb. 1992.
- [5] R. A. Sahner and K. S. Trivedi, "A hierarchical, combinatorial-Markov method of solving complex reliability models," in *Proc. ACM-IEEE Computer Society 1986 Fall Joint Computer Conference*, (Dallas, TX), pp. 817–825, Nov. 1986.
- [6] M. Veeraraghavan and K. S. Trivedi, "Hierarchical modeling for reliability and performance measures," in *Concurrent Computations* (S. K. Tewsburg, B. W. Dickinson, and S. C. Schwartz, eds.), pp. 449–474, Plenum Publishing Corporation, 1988.
- [7] M. Malhotra and K. S. Trivedi, "A methodology for formal expression of hierarchy in model solution," in *5th International Workshop on Petri Nets and Performance Models*, (Toulouse, France), pp. 258–267, Oct. 1993.
- [8] J. B. Dugan, K. S. Trivedi, M. K. Smotherman, and R. M. Geist, "The hybrid automated reliability predictor," *AIAA Journal of Guidance, Control and Dynamics*, vol. 9, no. 3, pp. 319–331, 1986.
- [9] R. Geist, "Extended behavioral decomposition for estimating ultrahigh reliability," *IEEE Trans. Reliability*, vol. R-40, pp. 22–28, Apr. 1991.
- [10] G. Ciardo and K. S. Trivedi, "A decomposition approach for stochastic reward net models," *Performance Evaluation*, vol. 18, no. 1, pp. 37–59, 1993.
- [11] A. T. Tai, K. S. Tso, W. H. Sanders, L. Alkalai, and S. N. Chau, "Low-cost flexible software fault tolerance for distributed computing," in *Proceedings of the 12th International Symposium on Software Reliability Engineering (ISSRE 2001)*, (Hong Kong, China), pp. 148–157, Nov. 2001.

- [12] B. Littlewood and D. Wright, "Stopping rules for the operational testing of safety-critical software," in *Digest of the 25th Annual International Symposium on Fault-Tolerant Computing*, (Pasadena, CA), pp. 444–453, June 1995.
- [13] A. T. Tai, K. S. Tso, L. Alkalai, S. N. Chau, and W. H. Sanders, "On the effectiveness of a message-driven confidence-driven protocol for guarded software upgrading," *Performance Evaluation*, vol. 44, pp. 211–236, Apr. 2001.
- [14] R. A. Sahner, K. S. Trivedi, and A. Puliafito, *Performance and Reliability Analysis of Computer Systems An Example-Based Approach Using the SHARPE Software Package*. Boston, MA: Kluwer Academic Publishers, 1995.
- [15] W. H. Sanders, W. D. Obal II, M. A. Qureshi, and F. K. Widjanarko, "The *UltraSAN* modeling environment," *Performance Evaluation*, vol. 24, no. 1, pp. 89–115, 1995.